

CLAIMS

1. A method of performing parity operations in a redundant data storage system that utilizes a plurality of data segments and at least two corresponding parity segments, wherein the parity segments are derived from the data segments and from parity coefficients corresponding to respective data segments, and wherein there are different parity operations involving different subsets of the parity coefficients, the method comprising the following steps:

pre-selecting parity coefficient subsets for use in the different parity operations;

storing all of the pre-selected parity coefficient subsets in a memory;

when performing a particular parity operation, determining which of the stored subsets of parity coefficients is needed for the particular parity operation;

reading the determined subset of parity coefficients from the memory;

performing the particular parity operation with the subset of parity coefficients that was read from the memory.

2. A method as recited in claim 1, wherein:

the storing step comprises pre-formatting the pre-selected parity coefficient subsets in an indexed memory array; and

the step of performing the particular parity operation is accomplished by hardware-based parity operation logic that utilizes the subsets of parity coefficients as they are pre-formatted in the memory array.

3. A method as recited in claim 1, further comprising:

classifying the different parity operations into classifications comprising:

1 parity segment generation operations;

2 parity segment regeneration operations;

3 data segment reconstruction operations;

4 wherein each classification of parity operations includes a plurality of
5 different classification scenarios, each classification scenario involving a
6 respective set of parity coefficients;

7 wherein the pre-selecting step comprises pre-selecting parity coefficient
8 subsets for each of the different classification scenarios;

9 wherein the storing step comprises pre-formatting the pre-selected parity
10 coefficient subsets in an indexed memory array; and

11 wherein the step of performing the particular parity operation is
12 accomplished by hardware-based parity operation logic that utilizes the subsets of
13 parity coefficients as they are pre-formatted in the indexed memory array.

14
15 4. A method as recited in claim 1, further comprising:

16 classifying the different parity operations into classifications comprising:

17 parity segment generation operations;

18 parity segment regeneration operations;

19 data segment reconstruction operations;

20 wherein each classification of parity operations includes a plurality of
21 different classification scenarios, each classification scenario involving a
22 respective subset of parity coefficients;

23 wherein the pre-selecting step comprises pre-selecting parity coefficient
24 subsets for each of the different classification scenarios;

1 wherein the storing step comprises grouping the pre-selected parity
2 coefficient subsets in an indexed memory array in accordance with the
3 classifications of the pre-selected parity coefficient subsets.

4
5 5. A method as recited in claim 1, further comprising:

6 classifying the different parity operations into classifications comprising:

7 parity segment generation operations;

8 parity segment regeneration operations;

9 data segment reconstruction operations;

10 wherein each classification of parity operations includes a plurality of
11 different classification scenarios, each classification scenario involving a
12 respective subset of parity coefficients, the subsets of parity coefficients having
13 varying subset sizes;

14 wherein the pre-selecting step comprises pre-selecting parity coefficient
15 subsets for each of the different classification scenarios;

16 grouping the pre-selected parity coefficient subsets into classification
17 groups;

18 sub-grouping at least one of the classification groups into sub-groups
19 according to subset size;

20 wherein the storing step comprises packing the pre-selected parity
21 coefficient subsets in an indexed memory array, segregated by group and sub-
22 group.

1 6. A method as recited in claim 1, wherein the different parity
2 operations include parity regeneration operations that include a plurality of
3 different operation scenarios, each operation scenario involving a respective subset
4 of parity coefficients, the respective subsets containing all possible ranges p_x
5 through p_y of a base set of N parameter coefficients referred to as p_0 through p_{N-1} ,
6 the method further comprising:

7 forming groups of the subsets according to subset size;

8 packing the groups of subsets in order of increasing subset size in a linear
9 memory array;

10 calculating an offset in the linear memory array of a particular group i
11 corresponding to a subset size L_i in accordance with a predefined function of N
12 and L_i ;

13 reading a subset of values from the particular group corresponding to the
14 subset size of L_i at the calculated offset in the linear memory array.

15
16 7. A method of accessing pre-selected subsets of values, wherein the
17 respective subsets contain all possible ranges p_x through p_y of a base set of N
18 values referred to as p_0 through p_{N-1} , the method comprising the following steps:

19 forming groups of the subsets, wherein the groups correspond respectively
20 to different subset sizes;

21 packing the groups of subsets in order of increasing subset size in a linear
22 memory array;

23 calculating an offset in the linear memory array of a particular group i
24 corresponding to a subset size L_i in accordance with a predefined function of N
25 and L_i ;

1 reading a subset of values from the particular group i at the calculated offset
2 in the linear memory array.

3
4 **8.** A method as recited in claim 7, wherein group i has $N - L_i + 1$
5 subsets of parity coefficients, and wherein the calculating step is performed by
6 evaluating a formula comprising:

$$((L_i - 1)(12N + (L_i)(3N - 2L_i - 5))/6) - 3(N - 1).$$

7
8
9 **9.** A method as recited in claim 7, further comprising:
10 within any individual group i , calculating a memory offset of a subset D in
11 the individual group as a function of subset size L_i .

12
13 **10.** A disk controller that performs parity operations in a redundant data
14 storage system that utilizes a plurality of data segments and at least two
15 corresponding parity segments, wherein the parity segments are derived from the
16 data segments and from parity coefficients corresponding to respective data
17 segments, and wherein there are different parity operations involving different
18 subsets of parity coefficients, comprising:

19 a memory containing stored subsets of parity coefficients corresponding
20 respectively to different parity-related computation scenarios;

21 hardware parity computation logic configured to any particular parity-
22 related operation by (a) determining the scenario of the operation. (b) reading the
23 corresponding subset of parity coefficients from the memory, and (c) performing
24 the particular parity-related operation with the subset of parity coefficients read
25 from the memory.

1
2 **11.** A disk controller as recited in claim 10, wherein the stored subsets
3 are indexed within the memory.

4
5 **12.** A disk controller as recited in claim 10, wherein:
6 the parity-related computation scenarios are classified under classifications
7 comprising:

8 parity segment generation operations;
9 parity segment regeneration operations;
10 data segment reconstruction operations;

11 the stored subsets are grouped in the memory in accordance with the
12 classifications of the corresponding parity-related computation scenarios.

13
14 **13.** A disk controller as recited in claim 10, wherein:
15 the parity-related computation scenarios are classified under classifications
16 comprising:

17 parity segment generation operations;
18 parity segment regeneration operations;
19 data segment reconstruction operations;

20 the stored subsets are grouped in the memory into classification groups, in
21 accordance with the classifications of the corresponding parity-related
22 computation scenarios.

23 at one of the classification groups is sub-grouped according to size of its
24 subsets.